



(12) **United States Patent**  
**Jacob et al.**

(10) **Patent No.:** **US 9,411,743 B2**  
(45) **Date of Patent:** **Aug. 9, 2016**

(54) **DETECTING MEMORY CORRUPTION**

(56) **References Cited**

(71) Applicant: **JUNIPER NETWORKS, INC.**,  
Sunnyvale, CA (US)

U.S. PATENT DOCUMENTS

(72) Inventors: **Samuel Jacob**, Bangalore (IN); **Vijay Paul**, Bangalore (IN)

5,418,956 A \* 5/1995 Willman ..... 711/206  
5,434,872 A \* 7/1995 Petersen ..... G06F 13/128  
370/445  
5,469,557 A \* 11/1995 Salt et al. .... 711/103  
5,974,500 A \* 10/1999 Maletsky ..... G06F 12/1433  
365/120  
6,009,495 A \* 12/1999 DeRoo et al. .... 711/103  
6,237,786 B1 5/2001 Ginter et al.  
7,228,404 B1 \* 6/2007 Patel et al. .... 712/228  
8,176,391 B2 \* 5/2012 Baysah ..... G06F 11/1044  
714/774  
8,621,337 B1 12/2013 Jacob et al.  
2002/0166036 A1 \* 11/2002 Watts, Jr. .... G06F 21/79  
711/163

(73) Assignee: **Juniper Networks, Inc.**, Sunnyvale, CA  
(US)

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 53 days.

(21) Appl. No.: **14/138,977**

(22) Filed: **Dec. 23, 2013**

(65) **Prior Publication Data**

US 2014/0122826 A1 May 1, 2014

**Related U.S. Application Data**

(63) Continuation of application No. 12/894,521, filed on  
Sep. 30, 2010, now Pat. No. 8,621,337.

(51) **Int. Cl.**

**G06F 11/00** (2006.01)  
**G06F 12/10** (2016.01)  
**G06F 11/07** (2006.01)

(52) **U.S. Cl.**

CPC ..... **G06F 12/10** (2013.01); **G06F 11/073**  
(2013.01); **G06F 11/0763** (2013.01)

(58) **Field of Classification Search**

CPC ..... G06F 11/0763; G06F 11/085; G06F  
11/0751; H05K 99/99; H04L 1/0061  
USPC ..... 714/811  
See application file for complete search history.

(Continued)  
OTHER PUBLICATIONS

“FreeBSD The Power to Serve,” [www.freebsd.org/cgi/man.cgi?query=memguard@sektion=9](http://www.freebsd.org/cgi/man.cgi?query=memguard@sektion=9), pp. 1-4, Aug. 2, 2010, FreeBSD 8.3.

*Primary Examiner* — Albert Decady

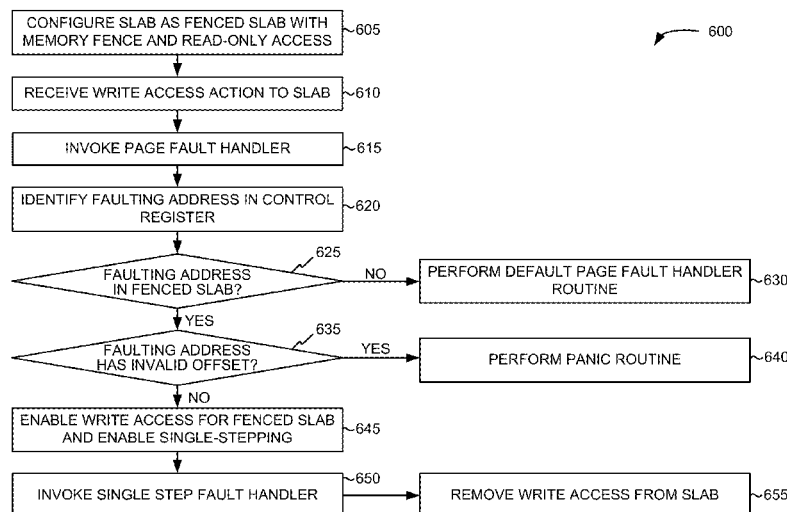
*Assistant Examiner* — Enam Ahmed

(74) *Attorney, Agent, or Firm* — Haritty & Haritty, LLP

(57) **ABSTRACT**

A device identifies, based on a program code instruction, an attempted write access operation to a fenced memory slab, where the fenced memory slab includes an alternating sequence of data buffers and guard buffers. The device assigns read-only protection to the fenced slab and invokes, based on the attempted write access operation, a page fault operation. When a faulting address of the attempted write operation is not an address for one of the multiple data buffers, the device performs a panic routine. When the faulting address of the attempted write operation is an address for one of the multiple data buffers, the device removes the read-only protection for the fenced slab and performs a single step processing routine for the program code instruction.

**20 Claims, 9 Drawing Sheets**



# US 9,411,743 B2

Page 2

(56)

## References Cited

### U.S. PATENT DOCUMENTS

2003/0228728	A1 *	12/2003	Yu .....	H01L 27/112	2008/0216073	A1 *	9/2008	Yates et al. ....	718/100
				438/201	2009/0100272	A1 *	4/2009	Smeets .....	713/189
2005/0120242	A1 *	6/2005	Mayer et al. ....	713/201	2011/0219708	A1 *	9/2011	Ohnishi et al. ....	52/79.1
					2012/0144167	A1 *	6/2012	Yates et al. ....	712/216

\* cited by examiner





Fig. 1



**Fig. 2**





Fig. 3

400 ↗



**Fig. 4**




Fig. 5

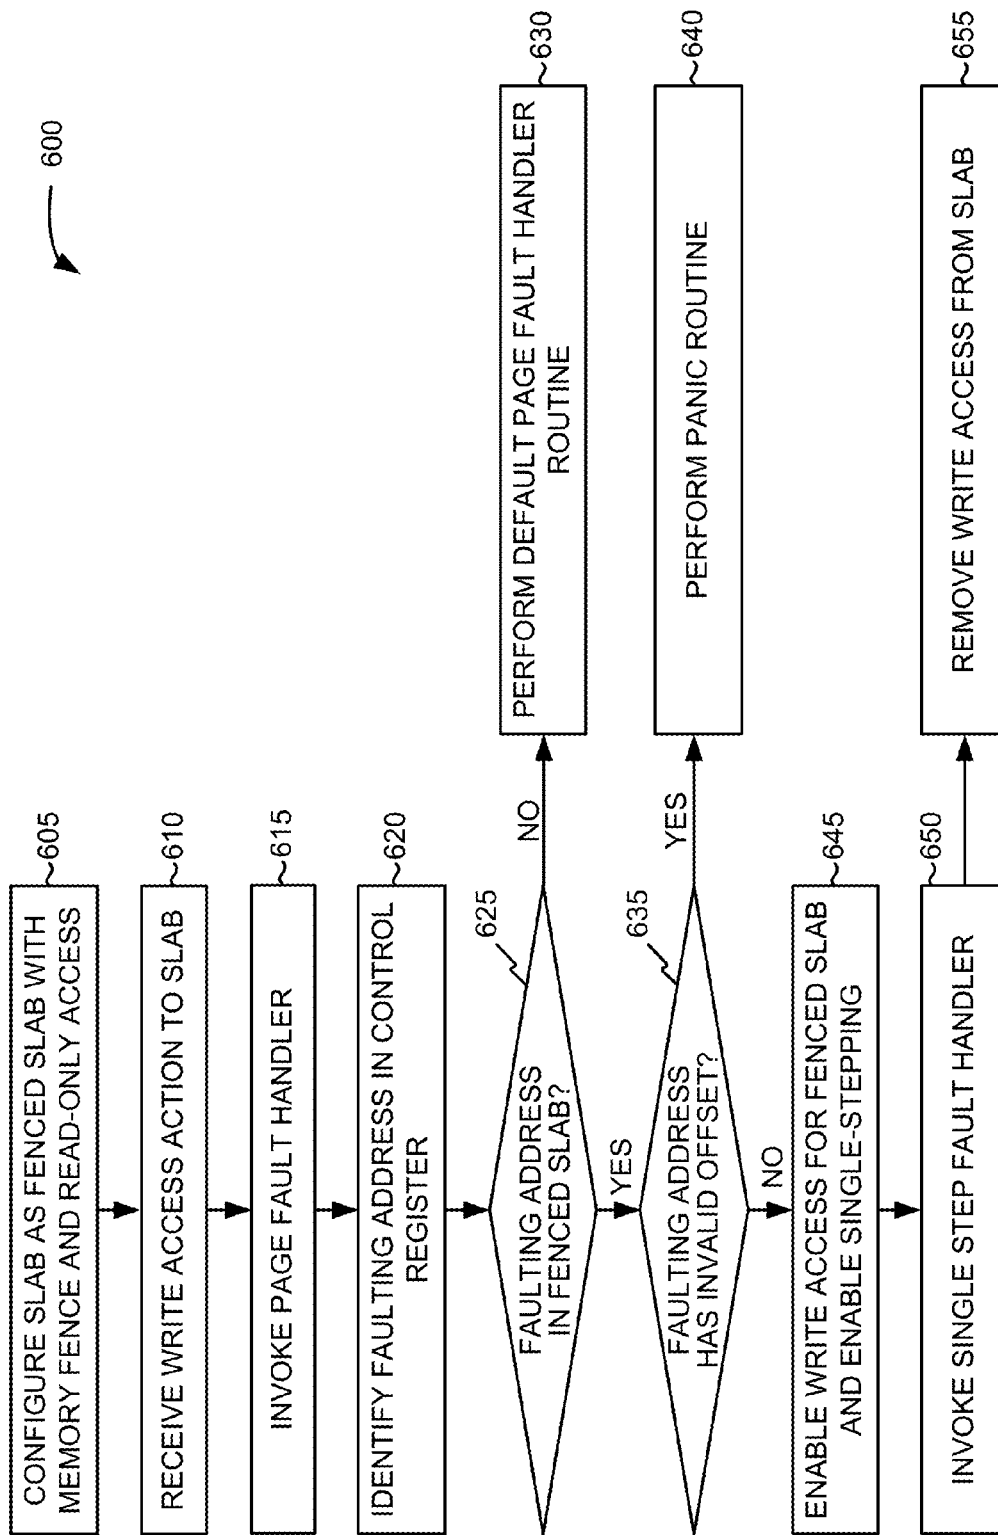


Fig. 6



700

```
static int  
Sysctl_test_fenc_alloc(SYSCTL_HANDLER_ARGS)  
{  
    int I;  
    char *va;  
  
    va = malloc(TEST_FENCE_SIZE, M_TEMP, M_WAITOK) ;  
    printf ("\n sysctl_test_fence_alloc() allocated VA %p\n", va) ;  
    for (i=0; i<TEST_FENCE_SIZE+1; i++)  
        va[i]=0;  
    return 0;  
}
```

PROBLEM  
CODE

Fig. 7

800

```

GEOM_LABEL: Label for provider ad0s1f is ufsid/4b6c7a4d499f5052.
947:(slab_zalloc)Fencing va 0xc3f75000 to c3f76000
947:(slab_zalloc)Fencing va 0xc3f76000 to c3f77000
947:(slab_zalloc)Fencing va 0xc3f77000 to c3f78000
947:(slab_zalloc)Fencing va 0xc3f78000 to c3f79000
947:(slab_zalloc)Fencing va 0xc3f7d000 to c3f7e000
947:(slab_zalloc)Fencing va 0xc3f7e000 to c3f7f000
947:(slab_zalloc)Fencing va 0xc3f7f000 to c3f78000
Trying to mount root from ufs:/dev/ad0s1a
WARNING: / was not properly dismounted
pcn0: link state changed to UP
Interface IP-Address Broadcast
Loading configuration files.
kernel dumps on /dev/ad0s1b
Entropy harvesting: interrupts ethernet_point to point
257:(handle_mem_fence_fault)VA 0xc3f7d080 IP 0xc07828c5 Page off 128
panic: FENCE: Access to protected item
cpuid = 0
KDB: enter: panic
[thread pid 73 tid 100044 ]
Stopped at    kdb_enter_why+0*3a: movl    $0,kdbwhy
db>
db>
db>

```

ACCESS TO  
INVALID ELEMENT  
CAUSES PANIC

Fig. 8

900

```

257: (handle_mem_fence_fault)VA 0xc3f7d080 IP 0xc07828c5 Page off 128
panic: FENCE: Access to protected item
cpuid = 0
KDB: enter: panic
[thread pid 73 tid 100044 ]
Stopped at   kdb_enter_why+0*3a: movl    $0,kdbwhy
db>
db>
db> bt
Tracing pid 73 tid 100044 td 0xc3f4d240
kdb_enter_why(c0825dc,c08256dc,c0840f79,c3b82a54,0,...) at kdb_enter_why+0*3a
panic(c0840f79,c3f7d080,c07828c5,80,2089aea0,...) at panic+0x136
trap(c3b82b04) at trap+0x1fd
calltrap() at calltrap+0x6
--- trap 0xc, eip = 0xc07828c5, esp = 0xc3b82b44, ebp = 0xc3b82b54 ---
sysctl_test_fence_alloc(c08b04c0,0,0,c3b82ba4,c3b82ba4,...) at sysctl_test_fence_
alloc+0x35
sysctl_root(c3b82ba4,c3b82bb0,4,c3b82bea,6400...) at sysctl_root+0x187
userland_sysctl(c3f4d240,c3b82c14,2,0,bfbf3434,...) at userland_sysctl+0x134
__sysctl(c3f4d240,c3b82cfc,18,6,c05d306b,...) at __sysctl+0x94
syscall(c3b82d38) at syscall+0x335
Xint0x80_syscall() at Xint0x80_syscall+0x20
--- syscall (202, FreeBSD ELF32, __syscall), eip = 0x2815b4cf, esp = 0xbfbfdbac, ebp =
0xbfbfdbb8 ---
db>

```

PROBLEM  
CODE  
PATH

Fig. 9

**DETECTING MEMORY CORRUPTION****RELATED APPLICATION**

This application is a continuation of U.S. patent application Ser. No. 12/894,521, filed Sep. 30, 2010 (now U.S. Pat. No. 8,621,337), the disclosure of which is incorporated herein by reference.

**BACKGROUND**

A computing device, such as a personal computer or a network device, may include an operating system (OS) that provides an interface between hardware of the computing device and software processes that are executed by the computing device. The OS may particularly include a kernel, which may be a central component of the OS and may act as a bridge between the software processes and the actual data processing done at the hardware level.

Memory corruption occurs when the contents of a memory location are unintentionally modified due to programming errors. When the corrupted memory contents are used later in a computer program, it may lead to a program crash or unexpected program behavior. Memory corruption in a kernel is difficult to debug because the corruption could be from any code (including third party drivers) and it could affect different data areas during each run. Some debugging tools attempt to detect memory corruption at a point in time after the corruption has occurred (e.g., at the time of free( ) in FreeBSD), at which point the corrupting code context is lost. Other tools can detect memory corruption at the exact moment of corruption, but require a lot of memory (usually a memory page even for 16 byte allocations). This large memory requirement poses problems when trying to debug problems in scaled scenarios where there is insufficient free memory to allow for usage of such tools.

**SUMMARY**

In one implementation, a method performed by a network device may include configuring, by a processor of the network device, a memory block as a fenced slab that includes multiple data buffers and multiple guard buffers, and assigning, by the processor, a read-only protection to the fenced slab. The method may further include identifying, by the processor, an attempted write access operation, based on a program code instruction, to one of the multiple data buffers or multiple guard buffers; recording, by the processor and in a control register, a faulting address for the attempted write access operation; invoking, by the processor and based on the identifying of the attempted write access operation, a page fault operation; retrieving, by the processor and from the control register, the faulting address; determining, by the processor, whether the faulting address is associated with the fenced slab; determining, by the processor and when the faulting address is associated with the fenced slab, whether the faulting address is an address for one of the multiple data buffers; performing, by the processor and when the faulting address is not an address for one of the multiple data buffers, a panic routine; removing, by the processor and when the faulting address is an address for one of the multiple data buffers, the read-only protection for the fenced slab; and performing, by the processor and when the faulting address is the address for one of the multiple data buffers, a single step processing routine for the program code instruction.

In another implementation, a device may include a memory to store instructions and a processor. The processor

may execute instructions in the memory to identify, based on a program code instruction, an attempted write access operation to a fenced memory slab, where the fenced memory slab includes an alternating sequence of data buffers and guard buffers; invoke, based on the attempted write access operation, a page fault operation; perform, when a faulting address of the attempted write operation is not an address for one of the multiple data buffers, a panic routine; and perform, when the faulting address of the attempted write operation is an address for one of the multiple data buffers, a single step processing routine for the program code instruction.

In a further implementation, a computer-readable memory having computer-executable instructions may include one or more instructions to identify, based on a program code instruction, an attempted write access operation to a fenced memory slab, where the fenced memory slab includes an alternating sequence of data buffers and guard buffers; one or more instructions to assign read-only protection to the fenced slab; one or more instructions to invoke, based on the attempted write access operation, a page fault operation; one or more instructions to perform, when a faulting address of the attempted write operation is not an address for one of the multiple data buffers, a panic routine; one or more instructions to remove, when the faulting address of the attempted write operation is an address for one of the multiple data buffers, the read-only protection for the fenced slab; one or more instructions to perform, when the faulting address of the attempted write operation is an address for one of the multiple data buffers, a single step processing routine for the program code instruction; and one or more instructions to re-assign, after the single step processing routine, the read-only protection to the fenced slab.

**BRIEF DESCRIPTION OF THE DRAWINGS**

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate one or more implementations described herein and, together with the description, explain these implementations. In the drawings:

FIG. 1 is a diagram of an example system in which concepts described herein may be implemented;

FIG. 2 is a block diagram of components of an example network device shown in FIG. 1;

FIG. 3 is a diagram of example components of a computing device shown in FIG. 1;

FIG. 4 is a diagram illustrating example functional components of the network device or the computing device of FIG. 1;

FIG. 5 is diagram illustrating an example data structure for a slab memory;

FIG. 6 is a flow diagram illustrating an example process for detecting memory corruption according to an implementation described herein; and

FIGS. 7-9 provide illustrations of memory corruption detection according to an implementation described herein.

**DETAILED DESCRIPTION**

The following detailed description refers to the accompanying drawings. The same reference numbers in different drawings may identify the same or similar elements. Also, the following detailed description does not limit the invention.

Systems and/or methods described herein may provide a debugging tool to detect memory corruption. The systems and/or methods may configure a read-only memory slab (or block) with an alternating sequence of data buffers and guard

buffers (referred to herein as a “fenced memory slab” or a “fenced slab”). When a program code instruction causes an attempted write access operation to any address within the fenced memory slab, a page fault operation may be invoked. If a faulting address of the attempted write operation is an invalid address (e.g., not an address for one of the multiple data buffers), a panic routine may be performed to identify the exact program code associated with the page fault. If the faulting address of the attempted write operation is a valid address (e.g., an address for one of the multiple data buffers within the fenced memory slab), the read-only protection for the fenced slab may be removed and a single step processing routine may be performed to provide a debug trace fault.

Implementations described herein are described primarily in the context of enhancing a slab memory allocator (or variations thereof) used in most of the operating system kernels based on x86 instruction set architectures (e.g., Intel 8086 instruction sets and successors). However, in other implementations, other memory allocators and/or architectures (e.g., a memory management unit (MMU)) can be used.

FIG. 1 is a diagram of an example system **100** in which concepts described herein may be implemented. System **100** may include a wide area network (WAN) **110** connected to one or more private networks **120-A** and **120-B** (collectively referred to as private networks **120**) and a computing device **130**. Private networks **120** may each, for example, include corporate or individual local area networks (LANs).

WAN **110** may generally include one or more types of networks. For instance, WAN **110** may include a cellular network, a satellite network, the Internet, or a combination of these (or other) networks that are used to transport data. Although shown as a single element in FIG. 1, WAN **110** may include a number of separate networks that function to provide services to private networks **120** and computing devices, such as computing device **130**. WAN **110** may be implemented using a number of network devices **115**. Network devices **115** may include, for example, routers, switches, gateways, and/or other devices that are used to implement WAN **110**.

Private networks **120** may each include a number of computing devices, such as, for example, client computing stations **125** and network devices **127**. Client computing stations **125** may include computing devices of end-users, such as desktop computers or laptop computers. Network devices **127**, similar to network devices **115**, may include network devices used to implement private networks **120**. For example, each of network devices **127** may include a data transfer device, such as a router, a gateway, a switch, a firewall, a network interface card (NIC), a hub, a bridge, a proxy server, an optical add-drop multiplexer (OADM), or some other type of device that processes and/or transfers traffic.

Network devices **115** and **127** may each implement a network operating system that controls the resources of the network device and provides an interface to the network device through which users can modify the configuration of the network device.

Computing device **130** may include, for example, a laptop or personal computer connected to WAN **110**. Alternatively, computing device **130** may include a mobile communication device, such as a cell phone, etc.

In the example system shown in FIG. 1, one WAN **110**, two private networks **120-A** and **120-B**, and one computing device **130** are shown. In other implementations, system **100** may include fewer devices, different devices, differently arranged devices, and/or additional devices than those depicted in FIG. 1. Alternatively, or additionally, one or more

devices of system **100** may perform one or more other tasks described as being performed by one or more other devices of system **100**.

FIG. 2 is a block diagram of an example network device **200**, which may correspond to one of network devices **115** or **127**. In order to increase throughput, network device **200** may use dedicated hardware to assist in processing incoming units of data, such as packets. In some alternative implementations, units of data (data units) other than packets may be used. As shown in FIG. 2, network device **200** may generally include a software portion **220** and a hardware portion **230**.

Software portion **220** may include software designed to control network device **200**. Software portion **220** may particularly include a network operating system (OS) **225**. For example, network operating system **225** may control hardware portion **230** and may provide an interface for user configuration of network device **200**. In general, software portion **220** may implement the functions of network device **200** that are not time critical. The functions described as being performed by software portion **220**, may be implemented through, for example, one or more general purpose processors **222** and one or more computer memories **224**. Processors **222** may include processors, microprocessors, or other types of processing logic that may interpret and execute instructions. Computer memories **224** (also referred to as computer-readable media herein) may include random access memories (RAMs), read-only memories (ROMs), and/or other types of dynamic or static storage devices that may store information and instructions for execution by one or more processors **222**.

Hardware portion **230** may include circuitry for efficiently processing packets received by network device **200**. Hardware portion **230** may include, for example, logic, such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), and/or a content-addressable memory (CAM). When network device **200** is a router, hardware portion **230** may, for example, receive incoming packets, extract header information for the packets, and process the packets based on the extracted header information. When network device **200** is a firewall, hardware portion **230** may, for example, receive incoming packets, extract header information from the packets, and match portions of the header information to a lookup table, such as one stored in a ternary CAM, to determine whether the packet should be dropped.

Network device **200** may additionally include one or more input ports **250** for receiving incoming packets and one or more output ports **255** for transmitting an outgoing packet. In some implementations, a port may act as both or one of an input port **250** or an output port **255**. Ports **250/255** may also be used to receive remote user connections for configuring the operation of network device **200**.

Although FIG. 2 shows example components of network device **200**, in other implementations, network device **200** may include fewer components, different components, differently arranged components, or additional components than depicted in FIG. 2. Alternatively, or additionally, one or more components of network device **200** may perform one or more other tasks described as being performed by one or more other components of network device **200**.

FIG. 3 is a diagram of example components of a computing device **300**, which may correspond to one of client computing stations **125** or computing device **130**. As shown in FIG. 3, computing device **300** may include a bus **310**, a processor **320**, a main memory **330**, a read only memory (ROM) **340**, a storage device **350**, an input device **360**, an output device **370**, and a communication interface **380**. Bus **310** may include a path that permits communication among the elements of the device.

5

Processor **320** may include a processor, microprocessor, or processing logic (e.g., an ASIC or a FPGA) that may interpret and execute instructions. Main memory **330** may include a RAM or another type of dynamic storage device that may store information and instructions for execution by processor **320**. ROM **340** may include a ROM device or another type of static storage device that may store static information and instructions for use by processor **320**. Storage device **350** may include a magnetic and/or optical recording medium and its corresponding drive.

Input device **360** may include a mechanism that permits an operator to input information to the device, such as a keyboard, a mouse, a pen, voice recognition and/or biometric mechanisms, etc. Output device **370** may include a mechanism that outputs information to the operator, including a display, a light emitting diode (LED), a speaker, etc. Communication interface **380** may include any transceiver-like mechanism that enables the device to communicate with other devices and/or systems.

As will be described in detail below, computing device **130** may perform certain operations in response to processor **320** executing software instructions contained in a computer-readable medium, such as main memory **330**. A computer-readable medium may be defined as a physical or logical memory device. A logical memory device may include memory space within a single physical memory device or spread across multiple physical memory devices. The software instructions may be read into main memory **330** from another computer-readable medium, such as data storage device **350**, or from another device via communication interface **380**. The software instructions contained in main memory **330** may cause processor **320** to perform processes that will be described later. Alternatively, hardwired circuitry may be used in place of or in combination with software instructions to implement processes described herein. Thus, implementations described herein are not limited to any specific combination of hardware circuitry and software.

Although FIG. **3** shows example components of computing device **300**, in other implementations, computing device **300** may include fewer components, different components, differently arranged components, or additional components than depicted in FIG. **3**. Alternatively, or additionally, one or more components of computing device **300** may perform one or more other tasks described as being performed by one or more other components of computing device **300**.

FIG. **4** is a block diagram of example functional components of a device **400** that may correspond to network device **115**, client computing device **125**, network device **127**, and/or computing device **130**. In one implementation, the functions described in connection with FIG. **4** may be performed by one or more components of network device **200** (FIG. **2**) or computing device **300** (FIG. **3**). As illustrated, device **400** may include a memory manager **410**, a page fault handler **420**, and a single step fault handler **430**. In one implementation, memory manager **410**, page fault handler **420**, and single step fault handler **430** may be included with an operating system (e.g., operating system **222**) of device **400**.

Memory manager **410** may include hardware or a combination of hardware and software to configure slab memory for device **400** and to detect invalid memory operations. In one implementation, slab memory may be configured as either an un-fenced (e.g., typical) slab or a fenced slab. Based on instructions from, for example, an operator, memory allocator **410** may assign, to a slab, appropriate values in a page table entry corresponding to a particular memory location. The page table entry may represent a physical address of a page in memory and a particular offset into the page. The page

6

table entry may also indicate permissions (e.g., read-only access, read/write access, etc.) for a particular slab. Slab memory configurations are described further in connection with FIG. **5**.

FIG. **5** provides a diagram of slab memory configurations according to an implementation described herein. As shown in FIG. **5**, slab memory component configurations may include an un-fenced slab **500** and a fenced slab **502**. In implementations described herein, a zone (e.g., a portion of memory) may be configured as one of un-fenced slab **500** or a fenced slab **502**. Fenced slab **502** may be configured, for example, to perform error detection for particular code.

Un-fenced slab **500** may represent a contiguous portion of memory (e.g. memory **224** or main memory **330**) that may include one or more physically contiguous pages. Un-fenced slab **500** may include multiple data buffers **510** and an on-page slab metadata portion **520**. Data buffers **510** may include space to store objects assigned by an OS kernel. On-page slab metadata portion **520** may include information needed to retain the slab, such as a slab header. In one implementation, each of data buffers **510** may be of equal size, and the number of data buffers may be sufficient to fill a one page (e.g., while leaving sufficient space for meta data portion **520**.) Un-fenced slab **500** may be configured with read/write access based on, for example, an entry in page table that corresponds to unfenced slab **500**.

Fenced slab **502** may represent another contiguous portion of memory (e.g. memory **224** or main memory **330**) that may include one or more physically contiguous pages. Fenced slab **502** may include multiple data buffers **510**, multiple guard buffers **530**, and on-page slab metadata portion **520**. More particularly, a guard buffer **530** may be placed after each data buffer **510**. Guard buffer **530** may include a portion of memory to which no data may be written. In one implementation, a size of guard buffer **530** may be equal to that of data buffer **510**. Fenced slab **502** may be protected from write access by setting the appropriate values in the page table entry corresponding to fenced slab **502**. Thus, whenever a kernel invokes a write access to fenced slab **502**, a page fault (or protection fault) routine may be invoked, regardless of whether the write access is actually valid or invalid.

Although FIG. **5** shows example arrangement of un-fenced slab **500** and fenced slab **502**, in other implementations, unfenced slab **500** and fenced slab **502** may be arranged differently than depicted in FIG. **5**.

Referring back to FIG. **4**, memory manager **410** may interrupt execution of a software program when the software program attempts to access pages that are not present in RAM or pages that do not conform to the permission attributes for a given page. For example, if a write operation is attempted for fenced slab **502**, memory manager **410** may invoke a page fault. When a page fault interruption occurs, memory manager **410** may store a page fault linear address in a control register (e.g., a CR2). Memory manager **410** may invoke page fault handler **420** to resolve the memory access problem.

Page fault handler **420** may include hardware or a combination of hardware and software to identify and resolve page faults received from memory manager **410**. In one implementation, page fault handler **420** may access the control register to identify a page fault linear address associated with the page fault. If the page fault linear address corresponds to one of un-fenced slabs **500**, page fault handler **420** may perform typical page fault operations to either resolve or report the page fault.

If the page fault linear address corresponds to one of fenced slabs **502**, page fault handler **420** may invoke steps to detect the point of the page fault. Page fault handler **420** may first

determine whether the write operation associated with the page fault linear address is within a correct offset. A correct offset may be considered an offset from the page which falls in one of data buffers **510** and not in metadata portion **520** or one of guard buffers **530**. Identification of an invalid offset may cause page fault handler **420** to invoke a panic routine. For example, if a write operation to one of guard buffers **530** occurs, page fault handler **420** may immediately default to a panic routine to identify corrupting code. The panic routine (or kernel panic) may generally output an error message to a console and/or dump an image of kernel memory to disk for purposes of debugging. The panic routine may wait for the system to be manually rebooted or initiate an automatic reboot.

If a write operation to one of buffers **510** (e.g., of fenced slab **502**) occurs, then page fault handler **420** may enable write access for fenced slab **502** (e.g., by modifying the appropriate values in the page table entry) and may invoke a single step processing routine. In one implementation, the single step processing routine may be performed by invoking single step fault handler **430**.

Single step fault handler **430** may include hardware or a combination of hardware and software to identify particular instructions that may have resulted in a particular page fault. For example, after write access is enabled for fenced slab **502**, the instruction of the application (e.g., the instruction that triggered the page fault) may be restarted by single step fault handler **430**. For example, single step fault handler **430** may step through the individual instruction(s) of the application program, may monitor how the processor (e.g., processor **222** or **320**) state changes after each instruction, and may identify errors based on the changes in the processor state. Thus, single step fault handler **430** may execute the application instruction without a page fault, but will generate a debug trace fault after executing the instruction. As part of the debug trace fault, single step fault handler **430** may remove write access (e.g., by modifying the appropriate values in the page table entry) for fenced slab **502** so that future write operations would again trigger a page fault.

Although FIG. **4** shows example functional components of device **400**, in other implementations, device **400** may include fewer functional components, different functional components, differently arranged functional components, or additional functional components than depicted in FIG. **4**. Alternatively, or additionally, one or more functional components of device **400** may perform one or more other tasks described as being performed by one or more other components of device **400**.

FIG. **6** is a flow diagram illustrating an example process **600** for detecting memory corruption in a kernel according to an implementation described herein. In one implementation, process **600** may be performed by device **400**. In another implementation, process **600** may be performed by another device or group of devices including or excluding device **400**.

Process **600** may include configuring a slab as fenced slab with a memory fence and read-only access (block **605**). For example, as described above in connection with FIGS. **4** and **5**, slab memory may be configured as either unfenced slab **500** or fenced slab **502**. Based on instructions from, for example, an operator, memory allocator **410** may assign, to a slab, appropriate values in a page table entry corresponding to a particular memory location. The page table entry may also indicate permissions (e.g., read-only access, read/write access, etc.) for a particular slab. Fenced slab **502** may represent a contiguous portion of memory (e.g. memory **224** or main memory **330**) that may include one or more physically contiguous pages. Fenced slab **502** may include multiple data

buffers **510** with a guard buffer **530** after each data buffer **510**. The size of guard buffer **530** may be equal to that of data buffer **510**. Fenced slab **502** may be protected from write access by setting the appropriate values in the page table entry corresponding to fenced slab **502**. Thus, whenever a kernel invokes a write access to fenced slab **502**, a page fault (or protection fault) routine may be invoked, regardless of whether the write access is actually valid or invalid.

A write action to the slab may be received (block **610**), a page fault handler may be invoked (block **615**), and a faulting address may be identified in a control register (block **620**). For example, as described above in connection with FIGS. **4** and **5**, fenced slab **502** may be protected from write access by setting the appropriate values in the page table entry corresponding to fenced slab **502**. Thus, whenever an application instruction invokes a write access to fenced slab **502**, a page fault (or protection fault) routine may be invoked, regardless of whether the write access is actually valid or invalid. Memory manager **410** may interrupt execution of a software program when the software program attempts a write operation for fenced slab **502**. When a page fault interruption occurs, memory manager **410** may store a page fault linear address in a control register. Memory manager **410** may then invoke page fault handler **420** to resolve the memory access problem.

Referring back to FIG. **6**, it may be determined if the faulting address is within the fenced slab (block **625**). If the faulting address is not within the fenced slab (block **625**—NO), a default page fault handler routine may be performed (block **630**). For example, as described above in connection with FIG. **4**, page fault handler **420** may access the control register to identify a page fault linear address associated with the page fault. If the page fault linear address corresponds to one of unfenced slabs **500**, page fault handler **420** may perform typical page fault operations to either resolve or report the page fault.

If the faulting address is within the fenced slab (block **625**—YES), it may be determined if the faulting address has an invalid offset (block **635**). If the faulting address has an invalid offset (block **635**—YES), a panic routine may be performed (block **640**). For example, as described above in connection with FIG. **4**, page fault handler **420** may determine whether a write operation associated with a page fault linear address is within a correct offset. Identification of an invalid offset may cause page fault handler **420** to invoke a panic routine. For example, if a write operation to one of guard buffers **530** occurs, page fault handler **420** may immediately default to a panic routine to identify corrupting code. The panic routine (or kernel panic) may generally output an error message to a console and/or dump an image of kernel memory to disk for purposes of debugging. The panic routine may wait for the system to be manually rebooted or initiate an automatic reboot.

If the faulting address does not have an invalid offset (block **635**—NO), write access for the fenced slab may be enabled and single-stepping may be enabled (block **645**), a single step fault handler may be invoked (block **650**), and write access may be removed from the fenced slab (block **655**). For example, as described above in connection with FIG. **4**, if a write operation to one of buffers **510** (e.g., of fenced slab **502**) occurs, then page fault handler **420** may enable write access for fenced slab **502** (e.g., by modifying the appropriate values in the page table entry) and may invoke a single step processing routine. In one implementation, the single step processing routine may be performed by invoking single step fault handler **430**. Single step fault handler **430** may restart the instruction that triggered the page fault. For example, single step

fault handler 430 may step through the individual instruction(s) of the application program, may monitor how the processor (e.g., processor 222 or 320) state changes after each instruction, and may identify errors based on the changes in the processor state. Thus, single step fault handler 430 may execute the application instruction without a page fault and may generate a debug trace fault after executing the instruction. As part of the debug trace fault, single step fault handler 430 may remove write access (e.g., by modifying the appropriate values in the page table entry) for fenced slab 502 so that future write operations would again trigger a page fault.

FIGS. 7-9 provide illustrations of memory corruption detection according to an implementation described herein. FIG. 7 provides an example code section 700 that includes instructions designed to cause memory corruption for a fenced slab. FIG. 8 is an example screen shot of a section 800 of a test output using code section 700. Section 800 shows that access for a first 128 elements (e.g. 0-127) was allowed and that access to an invalid element (128) caused a panic. Referring to FIG. 9, use of a backtrace ("bt") command reveals the particular problematic code path.

In the systems and/or methods described herein, a memory slab may be configured as a read-only fenced slab that includes multiple data buffers and multiple guard buffers. Write operations, based on program code instructions, to any address in the fenced slab may invoke a page fault. Once a page fault is invoked, the page fault handler may identify a particular faulting address and associate the faulting address with the fenced slab. If the faulting address is a valid address (e.g., for a data buffer), the page fault handler may temporarily remove read-only protection for the fenced slab and invoke a single step fault handler to execute the program code instruction. If the faulting address is an invalid address (e.g., for a guard buffer), the page fault handler may invoke a panic function.

The systems and/or methods described herein may provide for immediate detection of memory corruption and/or buffer overflow. The systems and/or methods described herein may be implemented without changes to existing memory alignment and access patterns, but may use less memory than existing techniques.

The foregoing description of example implementations provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention.

For example, while a series of blocks has been described with respect to FIG. 6, the order of the blocks may be varied in other implementations. Moreover, non-dependent blocks may be implemented in parallel.

It will be apparent that embodiments, as described herein, may be implemented in many different forms of software, firmware, and hardware in the implementations illustrated in the figures. The actual software code or specialized control hardware used to implement embodiments described herein is not limiting of the invention. Thus, the operation and behavior of the embodiments were described without reference to the specific software code—it being understood that software and control hardware may be designed to implement the embodiments based on the description herein.

Further, certain implementations described herein may be implemented as a "component" that performs one or more functions. This component may include hardware, such as a

processor, microprocessor, an application specific integrated circuit, or a field programmable gate array; or a combination of hardware and software.

It should be emphasized that the term "comprises" and/or "comprising" when used in this specification is taken to specify the presence of stated features, integers, steps, or components, but does not preclude the presence or addition of one or more other features, integers, steps, components, or groups thereof.

Even though particular combinations of features are recited in the claims and/or disclosed in the specification, these combinations are not intended to limit the disclosure of the invention. In fact, many of these features may be combined in ways not specifically recited in the claims and/or disclosed in the specification.

No element, act, or instruction used in the description of the present application should be construed as critical or essential to the invention unless explicitly described as such. Also, as used herein, the article "a" is intended to include one or more items. Where only one item is intended, the term "one" or similar language is used. Further, the phrase "based on," as used herein is intended to mean "based, at least in part, on" unless explicitly stated otherwise.

What is claimed is:

1. A method comprising:

determining, by a device, that a program code instruction causes an attempted write access operation to a memory block;

determining, by the device and after determining that the program code instruction causes the attempted write access operation, whether a faulting address, associated with the attempted write access operation, has an invalid offset;

removing, by the device, a read-only protection for the memory block when the faulting address does not have the invalid offset; and

performing, by the device and when the faulting address does not have the invalid offset, a single step processing routine to provide a debug trace for the program code instruction.

2. The method of claim 1, where determining whether the faulting address has the invalid offset comprises:

determining that the faulting address is an address for one of a plurality of memory buffers of the memory block, and

determining whether the faulting address has the invalid offset based on determining that the faulting address is the address for the one of the plurality of memory buffers.

3. The method of claim 1, further comprising:

configuring the memory block that includes a plurality of memory buffers and a plurality of guard buffers.

4. The method of claim 1, further comprising:

performing a panic routine when the faulting address is not valid.

5. The method of claim 4, where the panic routine identifies program code associated with a page fault.

6. The method of claim 1, where determining whether the faulting address has the invalid offset comprises:

determining that the faulting address is an address for one of a plurality of guard buffers in the memory block, and

determining whether the faulting address has the invalid offset based on determining that the faulting address is the address for the one of the plurality of guard buffers.



## 11

7. The method of claim 1, further comprising:  
invoking a page fault operation based on determining that  
the program code instruction causes the attempted write  
access operation to the memory block.
8. The method of claim 3, where each of the plurality of 5  
memory buffers is followed by one of a plurality of guard  
buffers.
9. A system comprising:  
one or more processors to:  
identify a faulting address associated with an attempted 10  
write access operation to a memory block;  
determine whether the faulting address, associated with  
the attempted write access operation, has an invalid  
offset;  
remove a read-only protection for the memory block 15  
when the faulting address does not have the invalid  
offset; and  
perform a single step processing routine when the fault-  
ing address does not have the invalid offset.
10. The system of claim 9, where the single step processing 20  
routine provides a debug trace for a program code instruction  
that causes the attempted write access operation.
11. The system of claim 9, where, when determining  
whether the faulting address has the invalid offset, the one or  
more processors are to: 25  
determine that the faulting address is within the memory  
block, and  
determine whether the faulting address has the invalid off-  
set based on determining that the faulting address is  
within the memory block. 30
12. The system of claim 9, where the one or more proces-  
sors are further to:  
invoke, when the faulting address is not valid, a panic  
routine to identify corrupting code.
13. The system of claim 12, where the panic routine outputs 35  
an error message to a console.
14. The system of claim 9, where the memory block is  
configured to include a plurality of memory buffers and a  
plurality of guard buffers.
15. The system of claim 9, 40  
where the one or more processors are further to:  
invoke a page fault operation based on the attempted  
write access operation, and

## 12

- where, when the single step processing routine is per-  
formed, the one or more processors are further to:  
remove write access for the memory block for a future  
write access operation to trigger a page fault.
16. A non-transitory computer-readable medium storing 5  
instructions, the instructions comprising:  
one or more instructions that, when executed by at least one  
processor, cause the at least one processor to:  
determine that a program code instruction causes an  
attempted write access operation to a memory block;  
determine whether a faulting address, associated with  
the attempted write access operation, has an invalid  
offset;  
remove a read-only protection for the memory block 15  
when the faulting address does not have the invalid  
offset; and  
invoke a processing routine to provide a debug trace for  
the program code instruction when the faulting  
address does not have the invalid offset.
17. The non-transitory computer-readable medium of  
claim 16, where the memory block is configured to include a  
plurality of memory buffers and a plurality of guard buffers.
18. The non-transitory computer-readable medium of  
claim 16, where the instructions further comprise:  
one or more instructions that, when executed by the at least 25  
one processor, cause the at least one processor to:  
determine that the faulting address is within a particular  
section of the memory block before determining  
whether the faulting address has the invalid offset.
19. The non-transitory computer-readable medium of  
claim 16, where the instructions further comprise:  
one or more instructions that, when executed by the at least  
one processor, cause the at least one processor to:  
remove write-access for the memory block after invok-  
ing the processing routine.
20. The non-transitory computer-readable medium of  
claim 16, where the instructions further comprise:  
one or more instructions that, when executed by the at least  
one processor, cause the at least one processor to:  
set one or more values in a page table entry, correspond-  
ing to the memory block, to provide the read-only  
protection for the memory block.

\* \* \* \* \*